

DETOx: Towards Optimal Software-based Soft-Error Detector Configurations

Michael Lenz and Horst Schirmeier
 Department of Computer Science 12
 Technische Universität Dortmund, Germany
 e-mail: {michael.lenz, horst.schirmeier}@tu-dortmund.de


Abstract—Application developers often place executable assertions – equipped with program-specific predicates – in their system, targeting programming errors. However, these detectors can detect data errors resulting from transient hardware faults in main memory as well. But while an assertion reduces silent data corruptions (SDCs) in the program state they check, they add runtime to the target program that increases the attack surface for the remaining state. This article outlines an approach to find an optimal subset of assertions that minimizes the SDC count, without the need to run fault-injection experiments for every possible assertion subset.

1. Introduction


With continuously shrinking semiconductor structure sizes and lower supply voltages, the per-device susceptibility to transient hardware faults is on the rise. A class of countermeasures with growing popularity is *software-implemented hardware fault tolerance* (SIHFT), which avoids expensive hardware mechanisms and can be applied application-specifically. However, SIHFT can, against intuition, cause more harm than good: its overhead in execution time and memory space also increases the system’s figurative “attack surface”. In consequence, this phenomenon can diminish all gains from detected or corrected errors by the increased possibility of being struck by radiation in the first place [1].

One class of SIHFT measures are *executable assertions* capable of detecting data errors. These statements or code sequences check statically-known state invariants in specifically chosen points of a workload at runtime, and “*can detect errors in input data and prevent error propagation*” [2]. Although primarily used as a means to complement unit tests – aiming at detecting programming errors during the development process especially of safety-critical software, – assertions can also detect data errors caused by hardware faults, or even be specifically designed for this purpose.

While, for example, Hiller et al. [3] identify good placements and adequate predicates for executable assertions, we assume the target workload is already equipped with a set of assertions. We explore the tradeoff between each assertion’s capability to detect data errors, and its runtime cost that increases the liveness of critical data in variables



(a) Assertions 1 and 2 both compiled in.



(b) Only assertion 2 compiled in: The runtime reduces by d_1 , the total gray area is reduced compared to Fig. 1a.

Figure 1: Fault space of an example workload in two assertion configurations: Gray areas denote SDCs, dark green areas show faults detected by assertion 2, and light green indicates detection by both assertions.

residing in memory during their execution. We also assume that data errors may be detected by more than one assertion, and aim at finding a subset of assertions – an *assertion configuration* – that catches most errors but minimizes the total *silent data corruption* (SDC) count.

2. The Attack-Surface Tradeoff

We assume a fault model of uniformly distributed, independent and transient single-bit flips in main memory, modeled as originating from *direct* influences of ionizing radiation [4]. Under this assumption, live data stored in a variable in memory has a corruption probability that is proportional to its lifetime during the workload run and its memory size [1] – or, in other words, proportional to its *area* in a fault-space diagram.

The X-axis of the example diagram in Fig. 1a shows the workload’s runtime, the Y-axis all bits in memory. In the example, *var2* holds live data throughout the complete

runtime, and a fault-injection (FI) campaign injecting faults at all possible time/memory-bit locations yields the information that bit-flips in this variable during this time frame lead to SDCs (color-coded *gray*). The gray-colored area contributes to the total SDC count of the workload.

Error Detection and Additional Runtime. Similarly, *var1* holds live data, but is checked against an application-specific predicate by executable assertions 1 and 2 at runtime. Hence, corruptions occurring in *var1* before the point in time when assertion 2 runs are *detected* (color-coded *green*). However, both assertions entail a runtime cost: The instructions executed for each assertion extend the workload runtime (by d_1 respectively d_2). Thereby, these assertions prevent corruptions in *var1* from causing an SDC, but increase the lifetime of the data stored in *var2* – which in turn enlarges the gray-colored area, or the total SDC count.

In consequence, depending on each assertion’s balance between SDC *reduction* (by turning otherwise gray SDC areas into green *detected* results) and SDC *increase* (by increasing the lifetime of other variables), it may influence the workload’s total SDC count positively or negatively. However, in the example in Fig. 1a, assertion 1 detects a subset of the data errors in *var1* that assertion 2 detects as well. Thus, it seems advisable to generate a workload variant *without* assertion 1, yielding a different fault-space diagram (Fig. 1b) with a lower total runtime and, hence, a lower (gray) SDC count originating from *var2*, but still most data errors in *var1* being detected by assertion 2.

Unfortunately, assertions cannot be considered independently: Depending on the assertion *configuration* – i.e., other assertions being compiled in or out of a specific workload variant – an assertion may increase the fault-space area of a protected (yielding green *detected* results) or an unprotected variable (gray SDC results). This interdependency makes it necessary to consider the total SDC count resulting from *every possible* assertion configuration.

Approach: FI-based Result Calculation. Considering real-world workloads with the number of N different executable assertions in the hundreds or thousands makes clear, that generating a workload variant for each of the 2^N possible assertion configurations is generally infeasible. Running an FI campaign measuring the total SDC count for each of these workload variants even exponentiates this problem.

Instead, our approach is based on FI results from a single workload variant with *all* assertions enabled. We exhaustively scan the single-bit flip fault space of this workload (using advanced pruning methods within the FAIL* FI tool [5]) and record a result – for example, *No Effect* or *SDC* – for each fault-space coordinate. If an FI experiment observes that one of the assertions detected a data error, we do not abort the experiment and record *detected* (green in Fig. 1), but record which assertion triggered and let the experiment continue running. So, for example, an FI experiment injecting in *var1* at some point in time before assertion 1 runs (Fig. 1a) would record assertion 1 to have detected the error, continue running, then also record assertion 2, and finally record that

an SDC occurred – because it *would* have occurred if none of the two assertions had been in place.

Based on this result data, our DETOx tool prototype can *calculate* the SDC count for an arbitrary assertion configuration. Removing, e.g., assertion 2 from the example in Fig. 1a requires 1) *subtracting* all gray SDC areas between $t_{2,1}$ and $t_{2,2}$ from the total SDC count, and 2) “*re-dyeing*” all remaining areas that were only detected by assertion 2 to the final result recorded in the FI experiments.

3. Sorting Example

Using our approach we evaluated the configurations of a simple sorting program taking 24 input elements with two assertions: the first one checks for ascending order of two swapped elements, while the second repeats the same test on the complete array after sorting. The following table shows the FI-obtained SDC count for the all-enabled (“11”) variant, and the predicted results for all other configurations. Assertion configurations are represented as a bit vector, where the bits indicate whether either assertion was enabled/disabled.

	00	01	10	11
SDC _{prediction}	2 315 851	1 395 495	2 318 461	1 547 549
SDC _{reality}	2 319 929	1 393 233	2 324 435	
error	-0.176 %	+0.162 %	-0.257 %	

To quantify the prediction quality, we ran FI campaigns for all other configurations besides “11” as a ground truth for comparison – information that would usually not be available, shown in the SDC_{reality} row. In this example, the SDC-count predictions are accurate to within 0.3 %. The optimal, lowest SDC-count configuration is “01”, i.e., only the second assertion gets enabled.

4. Conclusions

To conclude, our approach allows for fast and cheap exploration of the assertion-configuration space, and is based on FI results of a single, all-enabled configuration. This allows searching for optimal configurations in future work, for example using genetic algorithms or optimization techniques such as integer-linear programming.

References

- [1] H. Schirmeier, C. Borchert, and O. Spinczyk, “Avoiding pitfalls in fault-injection based comparison of program susceptibility to soft errors,” in *45th IEEE/IFIP Int. Conf. on Dep. Sys. & Netw. (DSN ’15)*. Washington, DC, USA: IEEE, Jun. 2015, pp. 319–330.
- [2] S. H. Saib, “Executable assertions – an aid to reliable software,” in *11th Asilomar Conf. on Circuits, Systems and Computers*. IEEE, 1977.
- [3] M. Hiller, A. Jhumka, and N. Suri, “On the placement of software mechanisms for detection of data errors,” in *32nd IEEE/IFIP Int. Conf. on Dep. Sys. & Netw. (DSN ’02)*. IEEE, Jun. 2002, pp. 135–144.
- [4] V. Sridharan, J. Stearley, N. DeBardeleben, S. Blanchard, and S. Gurumurthi, “Feng shui of supercomputer memory: Positional effects in DRAM and SRAM faults,” in *Int. Conf. for High Perf. Comp., Netw., Stor. & Analysis (SC ’13)*. New York, NY, USA: ACM, Nov. 2013.
- [5] H. Schirmeier, M. Hoffmann, C. Dietrich, M. Lenz, D. Lohmann, and O. Spinczyk, “FAIL*: An open and versatile fault-injection framework for the assessment of software-implemented hardware fault tolerance,” in *11th Europ. Depend. Comp. Conf. (EDCC ’15)*. IEEE, Sep. 2015.